arXiv:1610.04308v1 [cs.CV] 14 Oct 2016

Recurrent 3D Attentional Networks for End-to-End Active Object Recognition in Cluttered Scenes

Min Liu[1*], Yifei Shi[1*], Lintao Zheng[1] and Kai Xu[2]

***Abstract*— Active vision is inherently attention-driven: The agent selects views of observation to best approach the vision task while improving its internal representation of the scene being observed. Inspired by the recent success of attention-based models in 2D vision tasks based on single RGB images, we propose to address the multi-view depth-based active object recognition using attention mechanism, through developing an end-to-end recurrent 3D attentional network. The architecture comprises of a recurrent neural network (RNN), storing and updating an internal representation, and two levels of spatial transformer units, guiding two-level attentions. Our model, trained with a 3D shape database, is able to iteratively attend to the best views targeting an object of interest for recognizing it, and focus on the object in each view for removing the background clutter. To realize 3D view selection, we derive a 3D spatial transformer network which is differentiable for training with back-propagation, achieving must faster convergence than the reinforcement learning employed by most existing attention-based models. Experiments show that our method outperforms state-of-the-art methods in cluttered scenes.**

I. INTRODUCTION

Active object recognition plays a central role in robot-operated autonomous scene understanding and object manipulation. The problem involves planning the views of the visual sensor of a robot to maximally increase the accuracy and confidence of object recognition. Recently, active recognition has gained much boosting [1], thanks to the fast development of 3D sensing techniques (e.g., depth cameras) and the proliferation of 3D shape repositories. Our work also adopts the 3D geometric data-driven approach, under the setting of 2.5D depth acquisition.

For view selection, most existing works lie in the paradigm of information theoretic view evaluation (e.g. [2], [3]). For example, from a set of candidates, the view maximizing the mutual information between observations and object classes is selected. Such methods often present two issues. First, to estimate the mutual information, unobserved views must be sampled and the corresponding data must be synthesized from a learned generative model, making the view estimation inefficient [1]. Second, the object recognition model is typically learned independent on the view planner, although the two are really coupled in an active recognition system [4].

∗ Min Liu and Yifei Shi are joint first authors. This work was supported by NSFC (61572507, 61532003 and 61622212).

[1]Min Liu, Yifei Shi and Lintao Zheng are PhD students of School of Computer, National University of Defense Technology, 410073 Changsha, P. R. China

[2]Kai Xu is an Assistant Professor of School of Computer, National University of Defense Technology, 410073 Changsha, P. R. China kevin.kai.xu@gmail.com

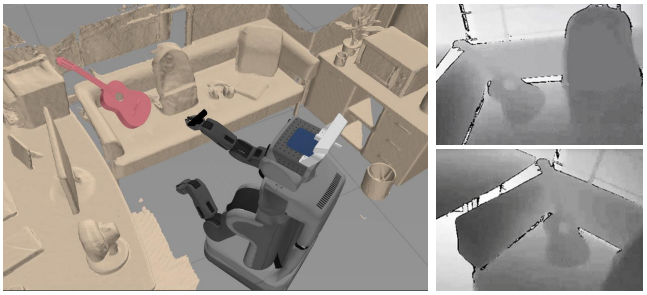

Fig. 1: Depth-based active object recognition in a cluttered indoor scene driven by a 3D attention-based view planning model. The best matching 3D model is retrieved from the 3D shape database and placed into the partially scanned scene.

Some works formulate active recognition as a reinforcement learning problem, to learn a viewing policy under various observations. Especially, a few recent works attempted end-to-end reinforcement learning based on recurrent neural networks [4], [5]. Applying a learned policy is apparently much more efficient than sampling from a generative model. However, these models are known to be hard to train, with difficult parameter tuning and relatively long training time [6]. Moreover, the success of these methods highly depends on the heuristically designed reward function.

The recent development of attention-based deep models has led to significant success in 2D vision tasks based on RGB images [6], [7]. Attention-based models achieve both efficiency and accuracy by focusing the processing only on the most informative parts of the input with respect to a given task. Information gained from different fixations is integrated into an internal representation, to approach the desired goal and guide future attention. Such mechanism, being both goal-directed and stimulus-driven [8], fits well to the problem setting of active recognition, accommodating object recognition and guided acquisition in a unified optimization.

However, the popular formulation of attention model based on recurrent neural networks [6] suffers from the problem of indifferentiable recognition loss over the attentional locations, making the network optimization by back-propagation infeasible. To make it learnable, the training is often turned into a partially observable Markov decision process (POMDP), which comes back to reinforcement learning. The recently introduced differentiable Spatial Transformer Networks (STN) can be used to actively predict image locations for 2D object detection and recognition [9]. Motivated by this, we opt to use STN units as our localization networks.

However, extending the standard STN to predict views in 3D space while keeping its differentiability is non-trivial.

To facilitate the back-propagation of loss gradient from a 2.5D depth image to 3D viewing parameters, we propose to parameterize the depth value at each pixel (x, y) in a depth image over the parameters of the corresponding view (θ, ϕ): $d(x, y) = f(\theta, \phi)$, through a ray-casting based depth computation at each pixel. Furthermore, to avoid the distraction of background clutter and make the recognition attend only to the foreground object of interest, we devise a second level 2D STN for foreground object localization in each view, saving the extra phase of static object detection [10].

Our two-level attention model produces efficient view planning and robust object recognition, as demonstrated by experimental evaluations on both synthetic and real data. We have integrated our method into the autonomous scene scanning system in [11], operated by a PR2 robot holding a depth camera, achieving active object recognition in a cluttered scene (Figure 1). Our main contributions include:

- A 3D attentional network for active object recognition through integrating RNN and STN.
- A differentiable extension of STN for view selection in 3D space, leading to a learnable attentional network which can be trained efficiently.
- A 2D localization network based on STN for localizing foreground object of interest in the presence of background clutter, achieving a robust recognition.

II. RELATED WORK

Active object recognition has a rich literature in robotics, vision and graphics (surveys available from e.g. [12], [13]). We provide a brief review of the related works (categorized into information theoretic and policy learning approaches), as well as some recent attempts on end-to-end learning.

Information theoretic approaches: Information theoretic formulation represents a standard approach to active vision problems. The basic idea is to quantify the information gain of each view by measuring the mutual information between observations and object classes [2], the entropy reduction of object hypotheses [14], or the decrease of belief uncertainty about the object that generated the observations [15]. The optimal views are those which are expected to receive the maximal information gain. The estimation of information gain usually involves learning a generative object model (likelihood or belief state) so that the posterior distribution of object class under different views can be estimated. Different methods have been utilized in estimating information gain, such as Monte Carlo sampling [2], Gaussian Process Regression [3] and reinforcement learning [16].

Policy learning approaches: Another line of research seeks to learn viewing policies. The problem is often viewed as a stochastic optimal control one and cast as a partially-observable Markov decision process. In [17], reinforcement learning is utilized to offline learn an approximate policy that maps a sequence of observations to a discriminative viewpoint. Kurniawati et al. employ a point-based approximate solver to obtain a non-greedy policy offline [18]. In contrast to offline learning, Lauri et al. attempted to apply Monte Carlo tree search (MCTS) to obtain online active hypothesis testing policy for view selection [19]. Our method learns and compiles viewing policies into the hidden layers of RNN, leading to a high-capacity view planning model.

End-to-end learning approaches: The recent fast development of deep learning models has aroused the interest of end-to-end learning of active vision policies [20]. Malmir et al. use deep Q-learning to find the optimal policy for view selection from raw images [21]. Our method shares similarities with the recent work of Wu et al. [1] in taking 2.5D depth images as input and 3D shapes as training data. They adopt a volumetric representation of 3D shapes and train a Convolutional Deep Belief Network (CDBN) to model the joint distribution over volume occupancy and shape category. By sampling from the distribution, shape completion can be performed based on observed depth images, over which virtual scanning is conducted to estimate the information gain of a view. Different from their method, our attention model is trained offline hence no online sampling is required, making it efficient for online active recognition. The works of Jayaraman and Grauman [4] and Xu et al. [5] are the most related to ours, where the recognition and control modules are jointly optimized based on reinforcement learning. We employ the spatial transformer units [9] as our locator networks to obtain a fully differentiable network.

III. APPROACH

We first provide an architectural overview of our recurrent attentional model, followed by detailed introduction of the individual subnetworks. Specifically, we elaborate the two levels of spatial transformer networks, responsible for view selection and foreground detection, respectively. Finally, we provide details about model training and inference.

A. Architecture Overview

Figure 2 shows the architecture of our recurrent attentional model. The main body of the model is a Recurrent Neural Network (RNN) for modeling the sequential dependencies among consecutive views. At each time step, the model takes a depth image as input, extracts image feature, amalgamates information of past views, makes a prediction of the object class and produces a next-best-view for future observation. To achieve that, we embed three subnetworks into the RNN: a 3D-STN for selecting views in 3D space, a 2D-STN for detecting object of interest in each view, and a shape classifier for depth-based object recognition.

Our model works as follows. Given an input depth image of the current view, it first selects within the image a rectangle region (not necessarily axis-aligned) containing the object being recognized, through estimating the transformation (including translation, rotation and scaling) of the region. This is done by a 2D localization network which is essentially a stack of convolutional layers (Conv_{2D}). The region is then transformed back to the input image size by the depth image generator (DG_{2D}) and fed into another stack of convolutional layers (Conv_{3D}) for feature extraction. The extracted features are aggregated with those extracted from the past views, with the help of the RNN hidden layers. The

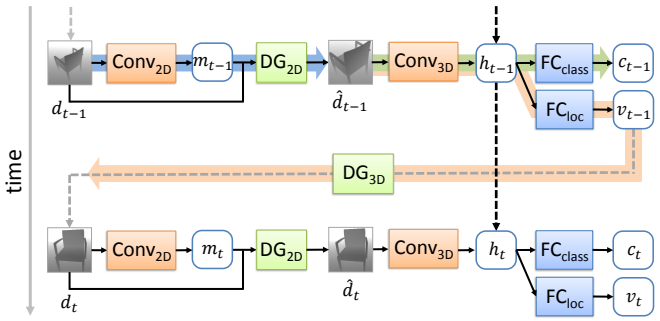

Fig. 2: Our recurrent attentional model. The dashed arrows indicate information flow across time steps while the solid ones represent that within a time step. The bold arrows underline the data flows of the three subnetworks, i.e., 2D-STN (blue), 3D-STN (orange) and CNN classifier (green).

aggregated features are then used for classification and next-best-view (NBV) prediction, with the fully connected layers FC_{class} and FC_{loc}, respectively. Based on the predicted NBV parameters, a depth image is generated (rendered or captured by the depth image generator DG_{3D}), serving as the input to the next time step.

As shown in Figure 2, the 2D-STN is the subnetwork constituted by the 2D localization network Conv_{2D} and the 2D depth image sampler DG_{2D}. The 3D-STN is the subnetwork encompassing the convolutional layers Conv_{3D}, the fully connected layers FC_{loc} and the depth image generator DG_{3D}. The shape classifier is composed of Conv_{3D} and FC_{class}, which is a standard Convolutional Neural Network (CNN) classifier. Therefore, the convolutional layers are shared by the CNN classifier and the 3D-STN.

B. Spatial Transformer Networks

Spatial transformer is a differentiable sampling-based network, which gives neural networks the ability to actively spatially transform the input data, without any spatial supervision in training. Generally, spatial transformer is composed of a localization network and a generator. STN achieves spatial attention by first passing the input image into a localization network which regresses the transformation, and then generating a transformed image by the generator. The transformed image is deemed to be easier to recognize or classify, thus better approaching the desired task.

STN fits well to our problem setting. Due to its differentiability, it enables end-to-end training with back-propagation, making the network easier to learn. It is relatively straightforward to employ it for object localization in the image of a given view. However, when using it to predict views in 3D, we face the problem of indifferentiable pixel values over viewing parameters, which is addressed by our work.

C. 2D-STN for Object Detection

Object recognition in a real scene is especially challenging due to scene clutters. The recognition of an object of interest is usually distracted by background clutter. To address this issue, a natural choice is to conduct object detection to remove the background distraction. However, object detection itself is a non-trivial task.

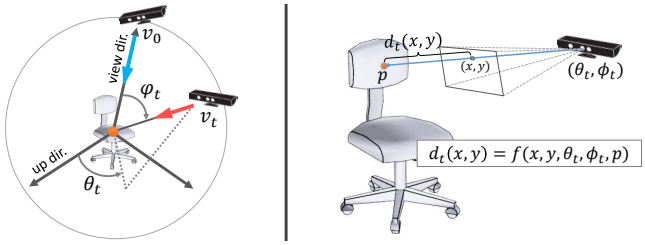

Fig. 3: Left: The viewing sphere around the object being recognized. The view of the current time step, v_t, is parameterized in the local spherical coordinate system built w.r.t. the initial view v_0. Right: Given a view $v_t = (\theta_t, \phi_t)$, the depth value of a pixel, $d_t(x, y)$, of the corresponding depth image is parameterized against the view parameters and the position of the intersection point, p, computed by ray-casting.

We approach object detection using spatial transformer networks due to its easy integration with our model. Since this STN performs localization in depth images, we refer to it as 2D-STN. There are two subnetworks in 2D-STN: a *2D localization network* Conv_{2D} and a *depth image generator* DG_{2D}. At a given time step t in the recurrent model (Figure 2), Conv_{2D} takes an input depth image (possibly with background) d_t and outputs the spatial transformation parameters m_t:

$$m_t = f^{\text{2D}}_{\text{Conv}}(d_t, W^{\text{2D}}_{\text{Conv}}), \tag{1}$$

where $W^{\text{2D}}_{\text{Conv}}$ are the weights to learn in Conv_{2D}.

The parameters are then used by DG_{2D} to generate a transformed image $\hat{x}_t$ containing only the foreground object:

$$\hat{d}_t = f^{\text{2D}}_{\text{DG}}(d_t, m_t). \tag{2}$$

Our 2D-STN differs from the original STN in [9] in a few aspects. First, for training data generation, unlike [9] where transformation and background noise are randomly added, we opt to use depth images of indoor scene as our background, to better fit our goal of object detection from an indoor scene. Specifically, we synthesize training data by using the depth of object as foreground while adopting the depth of scenes from the NYU dataset [22] as background. Since the background scenes also contain many household objects, 2D-STN can be enhanced by such training data, so that it is capable of detecting the foreground object being recognized, out of background clutter (see Figure 9 and 10). Second, the network is adapted to the size our input depth images (Figure 4). Third, since depth images are unlikely to distort with shear or anisotropic scaling, we remove the affine transformation in the standard STN.

D. 3D-STN for View Selection

Given an input depth image, the goal of 3D-STN is to extract image features, regress the 3D viewing parameters of the next-best-view (with respect to the recognition task) and generate (render or capture) a new depth image from that view. 3D-STN is again comprised of two subnetworks:

3D localization network (Conv_{3D} + FC_{loc}) and *depth image generator* DG_{3D}. During forward passing, the localization network takes the depth image d_t as input and outputs the next best view v_{t+1}. The viewing parameter is parameterized in the local spherical coordinate system built w.r.t. the initial view v_0 (Figure 3 left), and represented as a tuple (θ_t, φ_t), where θ_t is azimuth and φ_t elevation. Note that the viewing parameters does not include radial distance to the object center. This is because our 2D-STN is able to deal with object scaling more easily, thus simplifying the 3D-STN.

Specifically, the convolutional network Conv_{3D} first extract features from the depth image output by 2D-STN:

$$u_t = f^{\text{3D}}_{\text{Conv}}(\hat{d}_t, W^{\text{3D}}_{\text{Conv}}), \tag{3}$$

where $W^{\text{3D}}_{\text{Conv}}$ are the weights of Conv_{3D}. These features are amalgamated with those of past views stored in the RNN hidden layer h_{t-1}:

$$h_t = g(W_{ih}u_t + W_{hh}h_{t-1}), \tag{4}$$

where $g(\cdot)$ is a nonlinear activation function. W_{ih} and W_{hh} are the weights for input-to-hidden and hidden-to-hidden connections, respectively. The aggregated feature in h_t is then used to regress the viewing parameters of the next step:

$$(\theta_{t+1}, \varphi_{t+1}) = f^{\text{loc}}_{\text{FC}}(h_t, W^{\text{loc}}_{\text{FC}}), \tag{5}$$

where $W^{\text{loc}}_{\text{FC}}$ are the weights of FC_{loc}. The viewing parameters are then used by the depth image generator DG_{3D}, and to render a new depth image with the training 3D shape during training, or to capture one against the object of interest in testing:

$$d_{t+1} = f^{\text{3D}}_{\text{DG}}(\theta_{t+1}, \varphi_{t+1}). \tag{6}$$

In order to make our attentional network learnable, we require the generated depth image differentiable against the viewing parameters. This is achieved by casting a ray from the camera projection center passing through a pixel, and computing the intersection point between the ray and the 3D shape (if any). Therefore, as shown in Figure 3 (right), given the viewing parameters (θ_t, φ_t) and a 3D shape S, the depth value at any given pixel (x, y) of the depth image can be computed as: $d_t(x, y) = f_{\text{ray}}(x, y, p, \theta_t, \varphi_t)$, where p is the intersection point. f_{ray} simply computes the distance between p and pixel (x, y), which is differentiable. A basic assumption behind the parameterization of depth values w.r.t. viewing parameters and the intersection point is that the intersection point does not change when the view change is small, thus the first-order derivative over viewing parameters can be approximated by keeping the point fixed.

E. Classification Network

The depth image output by the 2D-STN at each time step is passed into a *CNN classifier* (Conv_{3D} + FC_{class}) for class label prediction. Note that the classification is also based on the aggregated features of both current and past views:

$$c_t = f^{\text{class}}_{\text{FC}}(h_t, W^{\text{class}}_{FC}), \tag{7}$$

where $W^{\text{class}}_{\text{FC}}$ are the weights of FC_{class}.

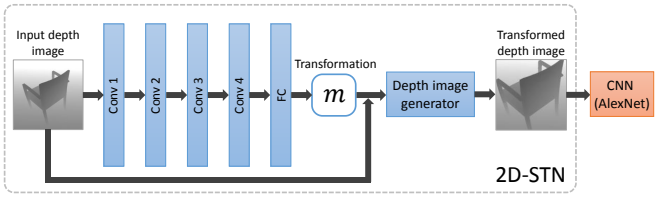

Fig. 4: The architecture of our 2D-STN (in the dashed box). The output of the network, which is the transformed depth image, is input into a CNN classifier (we use AlexNet [23]) for supervised training. Note that the AlexNet is not a part of our attentional model, but is pre-trained independently and used only for training the 2D-STN.

F. Training and inference

For the sake of convergence speed, we opt to decompose the training of our model into two parts: 1) 3D-STN, RNN and CNN classifier; 2) 2D-STN. The first part is trained by virtually recognizing the 3D shapes in the training dataset, using rendered depth images. For each 3D shape, we randomly select 50 views as the initial views. Starting from each initial view, we start the virtual recognizing for an episode of 10 time steps. In each step, the current depth image is input to the network and output both class label and the next-best-view. The network is then trained by two loss functions: the cross entropy loss in classification and the movement cost caused by the next-best-view. The movement cost is measured as the great circle arc length from the current view to the next-best-view in the spherical coordinate system. The training tunes the parameters of the 3D-STN, the RNN and the classification network simultaneously, using back-propagation through time (BPTT) [24].

To train the 2D-STN, we need to link it to a pre-trained CNN classifier (we use AlexNet [23]) as is shown in Figure 4, in order to train it with respect to an image classification task. During training, we fix the parameters of the AlexNet and only tune the parameters of the 2D-STN, thus making the training easier to converge.

At inference time, a captured depth image is passed into the our attentional model, for both object classification and next-best-view prediction. This is repeated until termination conditions are met. We set two termination conditions for inference: 1) The classification uncertainty, measured by the Shannon entropy of the classification probability, is less than a threshold (0.1); 2) The maximum number (15) of time steps has been reached.

IV. RESULTS & EVALUATION

A. Parameters and statistics

3D shape database for training: Our method has been trained and tested with a large-scale 3D shape database, i.e., the Princeton ModelNet40 dataset [1] containing 12,311 shapes in 40 object categories. 80% of the dataset was used for training and the remaining 20% for testing. During training, a 3D model needs to be rendered into a 224×224 depth image from any given viewpoint, using the ray-casting

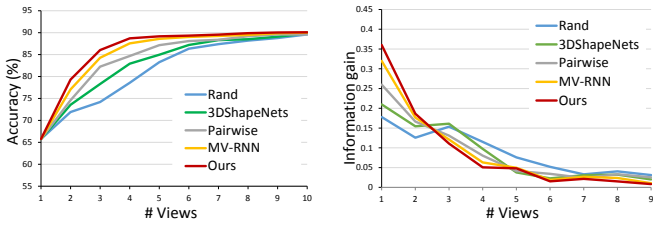

Fig. 5: Comparison of the performance of NBV selection by five methods (Random, 3DShapeNets [1], MV-RNN [27], Pairwise [25] and ours). Left: Recognition accuracy over the number of views. Right: Information gain of different views.

algorithm (parallel implementation on the GPU). All depth values are normalized into $[0, 255]$.

Parameters of subnetworks: The unit size for the hidden layer of the RNN is 1024. We use the ReLU activation functions for all hidden layers, $f(x) = max(0, x)$, for fast training. The classification network (CNN) takes the AlexNet architecture [23] which contains 5 convolutional layers and 3 fully connected layers followed by a soft-max layer. Same parameter settings as the original paper are used for the various layers. The 3D-STN contains 5 convolutional layers (shared with the classification CNN) and 2 fully connected layers. In summary, there are about $62M$ parameters in total being optimized in training the 3D-STN, the RNN and the CNN. The 2D-STN contains 4 convolutional layers containing respectively, as well as a fully connected layer. The first layer has 64 filters of size 11 and stride 4; the second layer has 192 filters of size 5; the second layer has 64 filters of size 3; the second layer has 20 filters of size 3. There are totally about $61M$ parameters in the 2D-STN.

B. Evaluation and comparison

Comparison of NBV estimation: To evaluate the performance of NBV estimation, we compare our attentional method against four alternatives, including a baseline method and three state-of-the-art ones. The baseline method *Random* selects the next view randomly. The state-of-the-art NBV techniques being compared include the *3DShapeNets* [1], the *MV-RNN* [11] and the active recognition method based on *Pairwise* learning [25]. The tests were conducted on the test shapes in ModelNet40 with the task of classification over the 40 classes. For a fair comparison on the quality of estimated NBVs, we let each method predict their own NBV sequences. and uniformly use a pre-trained Multi-View CNN shape classifier (MV-CNN) [26] for shape classification.

Figure 5(a) plots the recognition rate over the number of predicted NBVs. Our method achieves the fastest accuracy increase. To further evaluate the efficiency of the NBVs, we plot in (b) the information gain of different NBVs, measured by the decrease of Shannon entropy of the classification probability distributions. Compared to the alternatives, the NBV sequences output by our method attain larger information gain in early steps, demonstrating higher efficiency.

Timings: Table I lists the training and testing time of our method on both ModelNet10 (containing 4905 shapes in 10 categories) and ModelNet40 datasets. Since the 2D-STN is trained outside the main networks (RNN, CNN and 3D-STN), we report its training time separately. Table II compares the training time of three methods, i.e., 3DShapeNets, MV-RNN and ours. The training of 3DShapeNets involves learning the generative model with CDBN. MV-RNN is trained with reinforcement learning. The comparison shows the training efficiency of our model over the two alternatives. All timings were obtained on a workstation with an Intel® Xeon E5-2670 @ 2.30GHz × 24, 64GB RAM and an Nvidia® Quadro M5000 graphics card.

Dataset	Train		Test
	RNN + 3D-STN	2D-STN	
ModelNet10	7.0 hr.	2.7 hr.	0.1 sec.
ModelNet40	14.3 hr.	5.2 hr.	

TABLE I: Training and testing time of our method.

Method	Train	Test
3DShapeNets	96 hr.	4 min.
MV-RNN	66 hr.	0.1 sec.
Ours	19.5 hr.	0.1 sec.

TABLE II: Comparing training and testing time between 3DShapeNets, MV-RNN and our method.

NBV estimation with background clutter: The advantage of our method is best demonstrated in a cluttered scene where the recognition of a targeted object could be distracted by background clutter. With the 2D-STN, our method can robustly "detect" the object out of the background. We evaluate this on a *benchmark dataset* containing both synthetic and real test scenes, with inter-object occlusion (see Figure 7 for a few samples). Each targeted object in the test scenes was manually labeled with a ground-truth ModelNet40 category. We compare in Figure 6 the recognition accuracy of four methods including 3DShapeNets, MV-RNN and our method with and without the object localization network (2D-STN) on the two datasets, respectively. The comparisons show that our method, enhanced by the object localization in each view, increases the recognition accuracy by 8.9% and 10.3% on average for the synthetic and real datasets, respectively. Our method also significantly outperforms the two state-of-the-art methods being compared, demonstrating its strong resilience against background clutter.

Robustness against imperfect view: In real scenarios, a robot can hardly move the camera exactly on the viewing sphere around the object of interest, so the camera cannot always point to the center of the object while keeping a fixed distance to it. Such imperfect views may cause severe degrading of recognition accuracy, since the model had been trained with perfect viewing conditions. Here, a significant benefit of the 2D-STN is that our model can always "track" the object of interest under shifted viewing angle and varied viewing distance. To evaluate this, we test the robustness of our method against view perturbations. Specifically, we introduce view perturbations during the testing phase, through

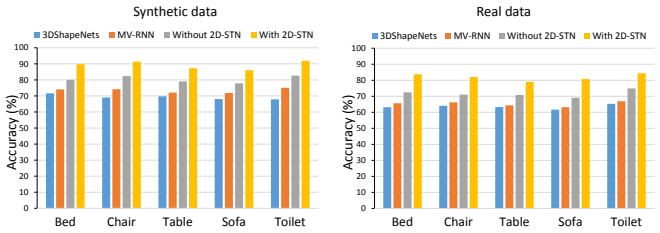

Fig. 6: Comparison of the accuracy of object recognition on scenes with background clutter (Fig. 7), using four methods: 3DShapeNets [1], MV-RNN [27], and our method with and without the object localization network (2D-STN).

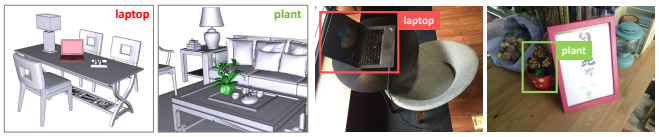

Fig. 7: A few samples of the synthetic (left) and real (right) test scene with background clutter.

altering the viewing direction of an estimated NBV by an angle up to $\pm 30^\circ$ and the viewing distance by up to $\pm 20\%$ of the original one. In Figure 8, we plot the recognition accuracy over the amount of view perturbation, for our method with and without the 2D-STN. The results demonstrate that our method gains robustness with the object localization, for perturbed views. For reference, the results of 3DShapeNets and MV-RNN are also given.

Visualization of attentions: To visually investigate the behavior of our attention model, we visualize in Figure 9 the sequence of NBVs along with the attended regions in the depth image at each view during the recognition of several real objects. We also test our algorithm on the newly released dataset of RGBD object scans [28], which provides RGBD videos targeting objects in various categories, captured by humans. Since there is no camera pose information, we run SLAM-based depth fusion [29] with each depth video and record the camera pose for each frame. Such parameterized depth video can serve as our viewing space, from which our method can "acquire" depth images, based on the viewing parameters. The results demonstrate that our method can correctly recognize the objects with plausibly planned views. After determining the object category, we utilize the MV-CNN method [26] to retrieve a 3D model (shown to the right) from the recognized category, using the ever-acquired depth images as input.

Testing on a robot platform: We have integrated our active object recognition with the autonomous scene scanning and reconstruction system propose in [11], to automate scene modeling via recognizing the objects in the scene, retrieving the best matched 3D shape and placing it into the scene, similar to [30] and [27]. The system is built on top of a Willow Garage PR2 robot holding a Microsoft Kinect (version 2), performing a rough scene reconstruction. Note that the reconstructed scene geometry was not used by our recognition, but only serves as the reference for scene modeling. Targeting to an object of interest, our method

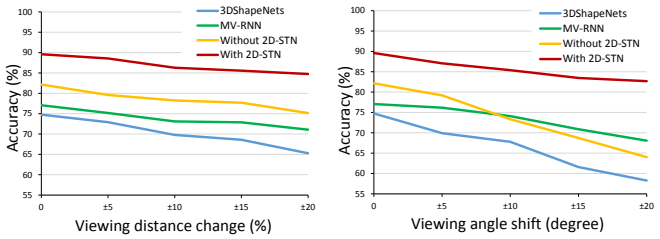

Fig. 8: Robustness against view perturbation. We plot the recognition accuracy against the amount of perturbation over viewing angle (left) and viewing distance (right). Our method with 2D-STN attains robustness against imperfect views.

drives the robot to acquire a series of depth images from the estimated NBVs, and output its category. We then retrieve a 3D model from the category as above, and place it into the scene through aligning the depth images against the scene geometry, similar to [27].

In Figure 10, we show the active recognition for scene modeling with the PR2 platform. In each depth image along the NBV sequence, we highlight the attended regions (corresponding to the object of interest). The final retrieved 3D model is placed into the partially reconstructed scene, shown to the right.

V. CONCLUSION

We have proposed a 3D attentional formulation to the active object recognition problem. This is mainly motivated by the mechanism resemblance between human attention and active perception [31], and the significant progress made in utilizing attentional models to address complicated vision tasks such as image captioning [7]. In developing such a model, we utilize RNNs for learning and storing the internal representation of the object being observed, CNNs for performing depth-based recognition, and STNs for selecting the next-best-views and detecting foreground object. The carefully designed 3D STN makes the whole network differentiable and hence easy to train. Experiments on both real and synthetic datasets demonstrate the efficiency and robustness of our active recognition model.

Limitations: A drawback of learning a policy offline is that physical restrictions during testing are hard to incorporate and when the environment is changed, the computed policy would no longer be useful. This problem can be alleviated by learning from a large amount of training data and cases using a high capacity learning model such as deep neural networks as we do. Although our method employs a 2D-STN for object localization, it does not handle mutual occlusion between objects which is a frequent case in cluttered scenes. One solution is to train the recognition network using depth images with synthetic occlusion. In our current model, the 2D-STN is trained separately. We have tested training it together with the main networks, however, it makes the training converge much slower but does not improve recognition performance notably.

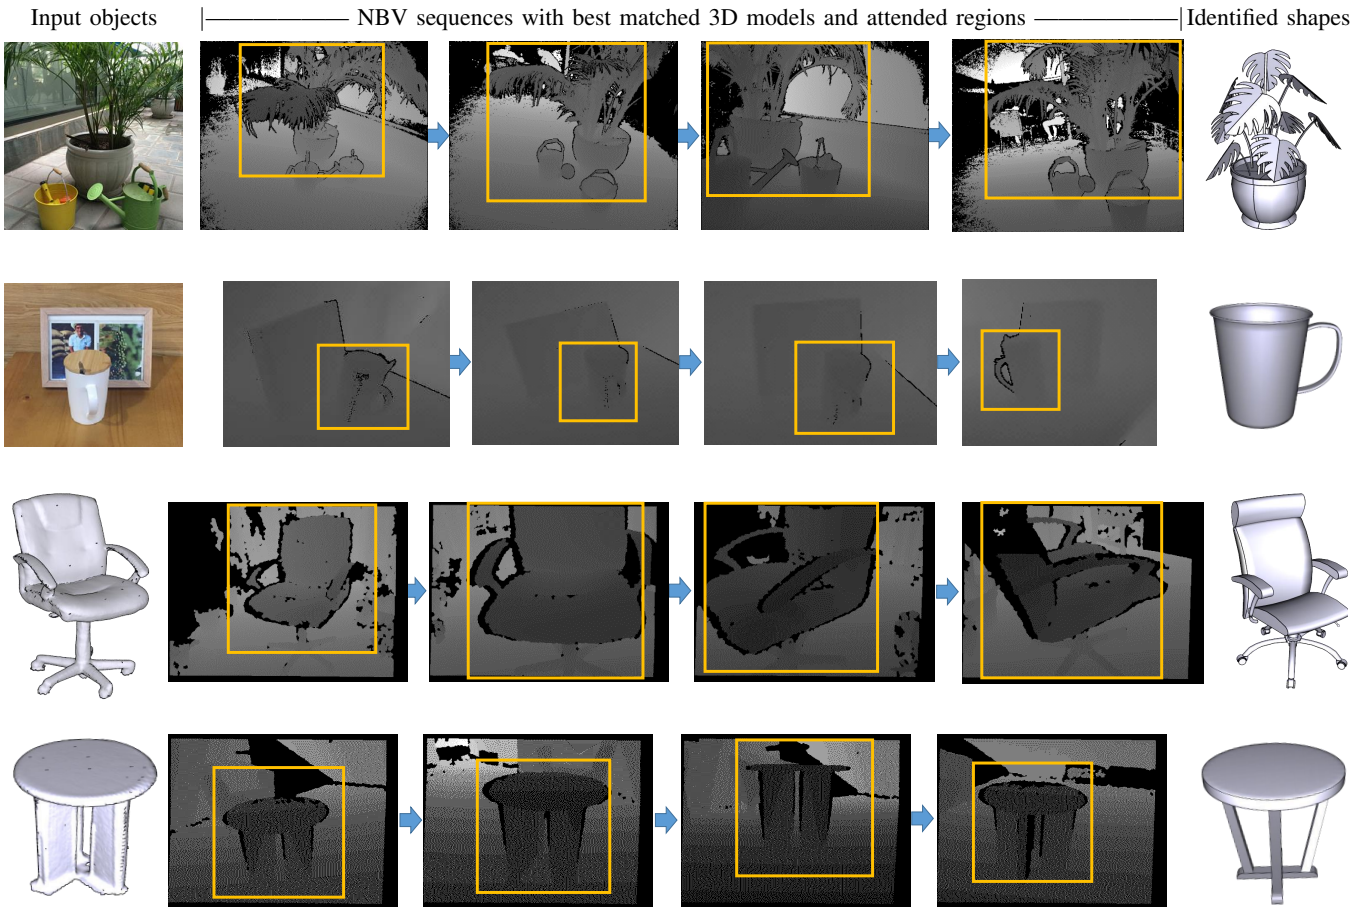

Fig. 9: Visualization of the active recognition process by showing the NBV sequence (depth images), the attended regions (orange boxes) in each depth image, running on two real scenes (top) and two scenes from the RGBD object scans database [28] (bottom). Note how our 2D-STN helps localize the object of interest in each view.

Future works: In the future, we would like to investigate a principled solution for handling object occlusion, e.g., using STN to help localize the shape parts which are both visible and discriminative, in a similar spirit to [32]. Another interesting direction is to study multi-agent attentions in achieving cooperative vision tasks, such as multi-robot scene reconstruction and understanding. It is particularly interesting to study the shared and distinct attentional patterns among heterogeneous robots such as mobile robots and drones.

REFERENCES

[1] Z. Wu, S. Song, A. Khosla, F. Yu, L. Zhang, X. Tang, and J. Xiao, "3D ShapeNets: A deep representation for volumetric shapes," in *Proc. CVPR*, 2015, pp. 1912–1920.

[2] J. Denzler and C. Brown, "An information theoretic approach to optimal sensor data selection for state estimation," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 24, no. 2, pp. 145–157, 2002.

[3] M. F. Huber, T. Dencker, M. Roschani, and J. Beyerer, "Bayesian active object recognition via gaussian process regression," in *International Conference on Information Fusion (FUSION)*, 2012, pp. 1718–1725.

[4] D. Jayaraman and K. Grauman, "Look-ahead before you leap: end-to-end active recognition by forecasting the effect of motion," in *Proc. ECCV*, 2016.

[5] K. Xu, Y. Shi, L. Zheng, J. Zhang, M. Liu, H. Huang, H. Su, D. Cohen-Or, and B. Chen, "3d attention-driven depth acquisition for object identification," *ACM Trans. on Graph. (SIGGRAPH Asia)*, vol. 35, no. 6, p. to appear, 2016.

[6] V. Mnih, N. Heess, A. Graves, *et al.*, "Recurrent models of visual attention," in *Proc. NIPS*, 2014, pp. 2204–2212.

[7] K. Xu, J. Ba, R. Kiros, A. Courville, R. Salakhutdinov, R. Zemel, and Y. Bengio, "Show, attend and tell: Neural image caption generation with visual attention," *arXiv preprint arXiv:1502.03044*, 2015.

[8] M. Corbetta and G. L. Shulman, "Control of goal-directed and stimulus-driven attention in the brain," *Nature Reviews Neuroscience*, vol. 3, pp. 201–215, 2002.

[9] M. Jaderberg, K. Simonyan, A. Zisserman, *et al.*, "Spatial transformer networks," in *Proc. NIPS*, 2015, pp. 2017–2025.

[10] N. Atanasov, B. Sankaran, J. L. Ny, G. J. Pappas, and K. Daniilidis, "Nonmyopic view planning for active object classification and pose estimation," *IEEE Trans. on Robotics*, vol. 30, no. 5, pp. 1078–1090, 2014.

[11] K. Xu, H. Huang, Y. Shi, H. Li, P. Long, J. Caichen, W. Sun, and B. Chen, "Autoscanning for coupled scene reconstruction and proactive object analysis," *ACM Trans. on Graph.*, vol. 34, no. 6, pp. 177:1–177:14, 2015.

[12] W. Scott, G. Roth, and J.-F. Rivest, "View planning for automated 3d object reconstruction and inspection," *ACM Computing Surveys*, vol. 35, no. 1, 2003.

[13] S. D. Roy, S. Chaudhury, and S. Banerjee, "Active recognition through next view planning: a survey," *Pattern Recognition*, vol. 37, no. 3, pp. 429–446, 2004.

[14] H. Borotschnig, L. Paletta, M. Prantl, and A. Pinz, "Appearance-based active object recognition," *Image and Vision Computing*, vol. 18, no. 9, pp. 715–727, 2000.

[15] F. G. Callari and F. P. Ferrie, "Active object recognition: Looking for differences," *International Journal of Computer Vision*, vol. 43, no. 3, pp. 189–204, 2001.

[16] T. Arbel and F. P. Ferrie, "Entropy-based gaze planning," *Image and vision computing*, vol. 19, no. 11, pp. 779–786, 2001.

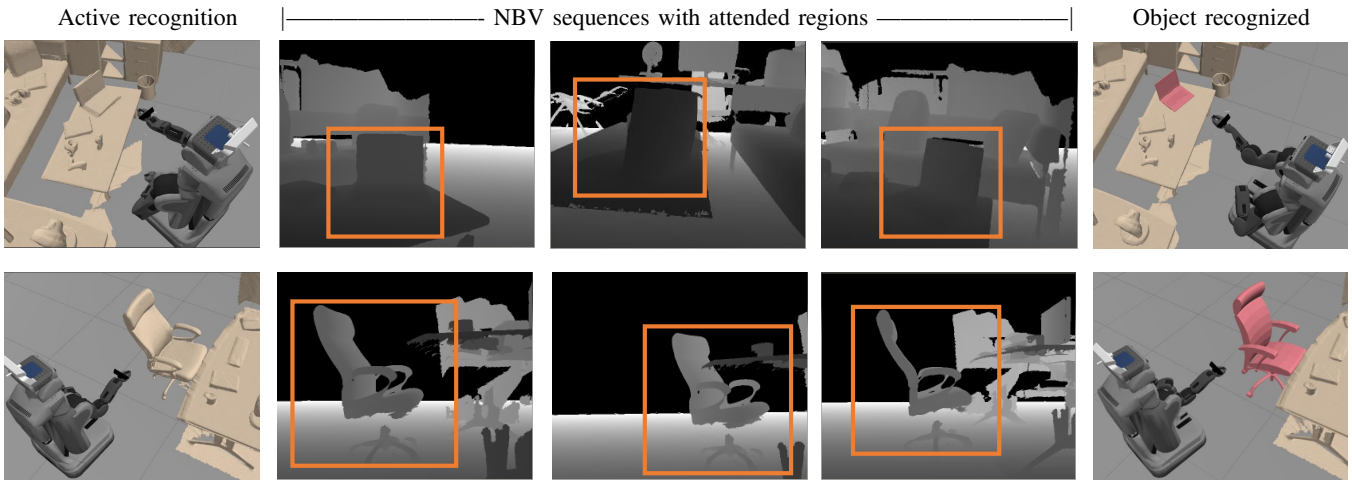

Fig. 10: Realizing our active object recognition on a PR2 robot platform. The recognition process is shown as the NBV sequence (depth images), along with the attended regions (orange boxes) in each depth image. The best matched 3D models, a laptop (top) and a chair (bottom), are retrieved and placed into the scenes (right).

[17] L. Paletta and A. Pinz, "Active object recognition by view integration and reinforcement learning," *Robotics and Autonomous Systems*, vol. 31, no. 1, pp. 71–86, 2000.

[18] H. Kurniawati, D. Hsu, and W. S. Lee, "Sarsop: Efficient point-based pomdp planning by approximating optimally reachable belief spaces." in *Robotics: Science and Systems*, vol. 2008, 2008.

[19] M. Lauri, N. Atanasov, G. Pappas, and R. Ritala, "Active object recognition via monte carlo tree search," in *Workshop on Beyond Geometric Constraints at the International Conference on Robotics and Automation (ICRA)*, 2015.

[20] S. Levine, C. Finn, T. Darrell, and P. Abbeel, "End-to-end training of deep visuomotor policies," *Journal of Machine Learning Research*, vol. 17, no. 39, pp. 1–40, 2016.

[21] M. Malmir, K. Sikka, D. Forster, J. Movellan, and G. W. Cottrell, "Deep q-learning for active recognition of germs: Baseline performance on a standardized dataset for active learning," in *Proceedings of the British Machine Vision Conference (BMVC)*, 2016, pp. 161–1.

[22] N. Silberman, P. Kohli, D. Hoiem, and R. Fergus, "Indoor segmentation and support inference from RGBD images," in *Proc. ECCV*, 2012, pp. 746–760.

[23] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Proc. NIPS*, 2012, pp. 1097–1105.

[24] M. C. Mozer, "A focused back-propagation algorithm for temporal pattern recognition," *Complex systems*, vol. 3, no. 4, pp. 349–381, 1989.

[25] E. Johns, S. Leutenegger, and A. J. Davison, "Pairwise decomposition of image sequences for active multi-view recognition," in *Proc. CVPR*, 2016.

[26] H. Su, S. Maji, E. Kalogerakis, and E. Learned-Miller, "Multi-view convolutional neural networks for 3D shape recognition," in *Proc. ICCV*, 2015.

[27] K. Xu, V. Kim, Q. Huang, and E. Kalogerakis, "Data-driven shape analysis and processing," *Computer Graphics Forum*, p. to appear, 2016.

[28] S. Choi, Q.-Y. Zhou, S. Miller, and V. Koltun, "A large dataset of object scans," *arXiv:1602.02481*, 2016.

[29] M. Nießner, M. Zollhöfer, S. Izadi, and M. Stamminger, "Real-time 3D reconstruction at scale using voxel hashing," *ACM Trans. on Graph. (SIGGRAPH Asia)*, vol. 32, no. 6, pp. 169:1–169:11, 2013.

[30] R. F. Salas-Moreno, R. A. Newcombe, H. Strasdat, P. H. J. Kelly, and A. J. Davison, "SLAM++: Simultaneous localisation and mapping at the level of objects," in *CVPR*, 2012, pp. 1352–1359.

[31] R. Bajcsy, "Active perception," *Proceedings of the IEEE*, vol. 76, no. 8, pp. 966–1005, 1988.

[32] T. Xiao, Y. Xu, K. Yang, J. Zhang, Y. Peng, and Z. Zhang, "The application of two-level attention models in deep convolutional neural network for fine-grained image classification," in *Proc. CVPR*, 2015, pp. 842–850.